RESEARCH PROPOSAL

COMPUTER LABORATORY FOR CLINICAL DECISION-MAKING

G.A. GORRY

MARCH 25, 1974

Introduction

As Schwartz has noted (1):

"Many discussions during the past decade have considered the use of computers as an adjunct to medicine. Few, however, have fully explored the possibility that the computer as an intellectual tool can reshape the present system of health care, fundamentally alter the role of the physician, and profoundly change the nature of medical manpower recruitment and medical education -- in short, the possibility that the health-care system by the year 2000 will be basically different from what it is today.

"Much has, of course, already been said about the role of the computer in improving the efficiency of the health-care system. These now familiar projections envision the computer performing a wide variety of functions such as the scheduling of hospital admissions, the keeping of medical records and the operation of laboratory and pharmacy. Such developments in the area of "house-keeping" activities offer considerable hope for the improvement of both hospital and outpatient operations but do not come to grips with the more fundamental problems of the health-care system -- the increasing shortage of physician manpower and the geographic maldistribution resulting from the reluctance of today's doctor to practice in rural or depressed urban communities. Even less do they give hope of dealing with the difficult challenge of maintaining a high level of physician competence in the face of a continued expansion of medical knowledge that tends to widen progressively the gap between what a doctor should know and what he can retain and utilize. The computer thus remains (in the light of conventional projections) as an adjunct to the present system, serving a palliative function but not really solving the major problems inherent in that system. There is, in fact, little reason to believe that any of the current proposals for solving these problems, technologic or other, will do more than mitigate their severity".

One radical and intriguing possibility for improving the efficiency and effectiveness of the health care system is to use the computer as an "intellectual" or "deductive" instrument -- a consultant that is built into the very structure of the health care system and augments the abilities of physicians and paramedical personnel. Clearly, however, considerable intellectual and technological resources must be marshalled and a long term research commitment must be made if this possibility is to be realized.

> ***We will argue in the body of this proposal, that the principal impediment to the realization of this exciting prospect is the lack of a good theory of clinical cognition. Despite successes in certain areas of clinical medicine, no theory of clinical decision-making has been developed which can explain the richness of the problem-solving behavior of***

> *experts. Further, we will argue that the computer is the key to the development of such a theory. The computer provides an environment in which ideas about* <u>*process*</u> *can be expressed in a quite natural way. Such environment is essential if we are to advance our understanding of clinical cognition.*

Although the idea of using computers and computer programs in the development of cognitive theories is not new (2), recent developments in computer science and technology make this idea more powerful. We have organized a team of computer scientists and medical scientists in a concerted attack on the problem of understanding clinical decision-making in new and profound ways. The Computer Laboratory concept is one which fits well into our current activities, and indeed, it offers us real leverage with respect to the growth of our efforts.

<u>What We Propose To Do</u>

Various approaches to the problems of automating processes for clinical decision-making have been employed by researchers in the field, and considerable success has been achieved. We believe that an expert program which can deliver advice and consultation with respect to serious clinical problems will make use of many of these approaches. At present, however, none of these approaches is sufficiently powerful to offer the <u>integrative</u> or <u>administrative</u> capability required to organize the variety of problem solving approaches necessary for the full range of clinical problems. Thus while other researchers continue with the development and refinement of existing techniques, we propose to devote our efforts to the problem of defining and implementing the framework within which these techniques can be organized and controlled.

The only examples we have of the integrative abilities which are required come from the performance of clinical experts. Clearly they possess the administrative problem-solving knowledge to shift from one approach to another as the case merits. For this reason, the principal focus of our efforts will be on gaining a better understanding of the behavior of experts.

We propose to undertake a program of research which will result in a new and significantly better theory of clinical cognition, with special emphasis on the administrative aspects of the problem-solving behavior. The computer will play a central role in the formulation and testing of this theory. Further, because the concepts upon which the theory will be based will be expressed in a form which is programmable, we will have a new technological framework within which efforts to create <u>distributable</u> <u>expertise</u> can proceed in concert. This in turn will speed the realization of that revolutionary role of the computer in the health-care system suggested above.

The activities of the Laboratory initially will be centered on several specific research projects which are related to our overall goal. These projects are discussed in detail in the body of this request. Here we will simply mention them and their relation to our primary goal.

1) Taking the Present Illness

The present illness is the initial point of contact between the patient and the physician, and for this reason, it represents a logical starting point. More importantly, however, the cognitive demands of taking the present illness, establishing the facts, drawing inferences about the facts and about the patient, dealing with discrepant information and uncertainty,. etc. are central to all clinical decision-making.

One of our major projects will be to develop a computer simulation of an expert taking the present illness. Such a simulation will be based on specific mechanisms for solving the various cognitive problems involved. These mechanisms, in turn, will be central to a variety of other decision-making programs. The knowledge gained from this effort and the results of the next project discussed will allow us to attack the problems of differential diagnosis and the risk/benefit analysis of management.

2) The Formalization of Clinical Knowledge

A second major project of the Laboratory will involve the development of new ways to formalize medical knowledge. Initially, this knowledge will be primarily that which appears in texts or journal articles on clinical problems, augmented and refined by clinical experts.

The criteria by which proposed representations of this knowledge will be judged include:

a) clarity
b) parsimony
c) completeness
d) capacity for expressing relations among "pieces" of knowledge
e) the ease with which it can be assimilated by a computer

Loosely speaking, the present illness project can be said to be concerned with how knowledge is used, whereas this project is concerned with formalizing what knowledge is required.

The result of this effort will be a methodology for building a knowledge base for programs such as the cognitive simulation of the present illness, and it should be viewed as being intimately connected with that project. Further it should provide a basis for research on the construction of diagnostic and management programs for various problems by providing a framework within which the basic knowledge required can be organized. The Laboratory will also develop programs for diagnosis and management when a good base of understanding has been achieved.

3) Model-Based Decision-Making

There are a number of important areas of clinical medicine in which a formal (generally mathematical) model is available upon which certain diagnostic or management decisions could conceivably be based. In many of these cases, however, the model in question is of little clinical use. Although the model often surpasses the ability of even the best physician to deal with certain aspects of the problem, or with "classic" cases, it cannot cope with a variety of patient-specific factors which should be factored into the decisions, or certain emergency conditions which should cause a re-ordering of the priorities in the model. In general physicians understand how to alter and refine their approach to a problem in the light of such factors, but computer programs unfortunately remain very rigid in this regard.

If we look to the day when various models and techniques are combined in a single system, it is clear the new flexibility must be built into the component pieces so that they can be "tailored" to fit a certain situation, and so that component pieces work coherently under the same assumptions about the patient.

To achieve this aim, we need new ways to combine medical "common sense" with mathematical models. The models themselves must be represented in such a way as to allow this common sense to be applied. Hence it must be clear to some supervisory program what the basis for a particular model is, and how changes in assumptions about the patient affect this basis, and hence the model.

We will begin to investigate these problems in the context of a model for the administration of digitalis/digoxin. This problem is a good one, because the "best" strategy for any patient depends in part on the use of a model, and in part on a basic understanding the medical problems of the patient.

Some Recurrent Problems

There are several problems which arise in almost all phases of clinical decision-making, and these will be the focus of a continuing research activity of the Laboratory. We mention them separately here, but we want to emphasize that they really represent threads which run through all our work.

4) Dealing with Discrepant Information

One of the important problems in clinical medicine is the amount of discrepant information which must be dealt with. Some of this difficulty arises because patients are not always accurate observers of their symptoms, or because they wish to conceal facts from the physician. Other problems arise from errors in laboratory tests or medical records. In addition there are many problems in which the discrepancy is not absolute, but rather relative to some currently believed hypothesis about the patient.

The question of belief is thus central to clinical decision-making. We plan to study this problem in a variety of contexts, with the intention of answering such questions as:

> How is the credibility of a piece of information established?
> How are potential discrepancies among facts detected?
> How are conflicts between facts resolved?
> What strategies are employed to resolve ambiguities or discrepancies?

5) The Representation of Time

Time plays a key role in clinical medicine. Diseases and their manifestations evolve through time. The interpretation of facts is often affected by the place of these facts in time. Often time-based relationships are crucial in making diagnoses or management decisions.

If we are to capture clinical expertise in a machine, we must equip the machine with an understanding of time and events which take place in time. Thus the machine needs a minimal ability to place events and intervals on some form of "time-line", and to make appropriate deductions about this arrangement. But much more is required. For example, we must develop ways to capture the concept of episodes. The machine needs to understand such fragments as "the gradual onset of the disease" and "an abrupt cessation of symptoms".

This is an area where substantive progress can probably be of direct use to other researchers in the field who to date have employed rather ad hoc methods to solve the problems of time representation or who have had to skirt the issue entirely to the detriment of their efforts.

6) <u>Inquiry and Explanation</u>

Another area in which we will be working is the development of mechanisms which allow a user to employ a natural and direct mode of interaction with a program and will allow a program the ability to explain its behavior in terms which are readily understandable to a clinician.

To a large extent, we will rely on research and development of natural language capabilities by others, in particular some of our colleagues at M.I.T, but we will play an active role in adapting their work to the medical context.

We will play a more central role in the development of the technology which will allow a program to generate explanations. Such explanations may be based on a variety of principles such as the use of physiological models. The point is that such a capability must be developed to meet several needs:

a) the need for users to understand the basis for a program's advice, particularly when the clinical problem is a serious one.
b) the need for clinicians working in our group to have access to facts and procedures used by the program in arriving at a particular conclusion.
c) the need for students to interrogate the program to learn about its strategies

Here again, progress in the development of these facilities, coupled with progress on our other projects should have an immediate and direct impact on the work of other researchers in the field, as well as a longer term impact on the delivery of health care.

<u>Summary</u>

In summary, we are proposing some projects which we believe will provide the proper direction for the Laboratory. The problems addressed by these projects are all basic problems for computer-aided clinical decision-making. Our emphasis on the study of clinical experts and on the use of the latest concepts of computer science to express the results of this study will provide a unifying theme for members of the Laboratory.

We have already formed a group of computer scientists, clinicians, and graduate students, which has begun work on these problems. The Laboratory would greatly facilitate and accelerate collaborative efforts of this kind, and it would be a link between the impressive computer science resources of M.I.T. and the equally impressive clinical resources of the Tufts-New England Medical Center. It would also provide a center into which researchers from other institutions could be drawn. In all, we envision that the Laboratory would be the center of new, vital, and important combinations of research and education. Its activities should have a significant impact on the computer-aided delivery of health care, as well as on medical education.

Background

The Laboratory we are proposing here will bring together experts from the computer sciences and from medicine for the purpose of gaining a new and deep understanding of the processes of clinical cognition and developing the mechanisms to translate this understanding into improvements in health care delivery. Here we want to give a brief history of the development of the research group, and then because of our involvment in both medicine and computer science, we want to briefly review important concepts and developments in both computer-aided clinical decision-making and in the relation of computer science to psychology and to theories of problem-solving.

The Development of Our Research Group

In order to put our application into perspective, we want to include a brief history of the development of the research group.

The nucleus of the group was formed several years ago, and it consisted of Drs. Schwartz and Kassirer and Professor Gorry. Schwartz and Kassirer had been working on the problem of encoding the protocols of experts in computer programs, and had developed a program for acid-base problems {1}. Gorry had developed a program which used statistical decision theory to solve diagnostic problems {3}. Because of the common interest in automating processes for clinical decision-making, the three joined forces.

The initial efforts of the group were directed along the lines suggested by the decision theory program. The work was considerably deepened and expanded during the two years following the initial formation of the group. A series of papers describing the work were published, most notably two recent articles ({4 } and {5}). These two papers consider in detail the application of decision analysis to clinical decision making, both insofar as the automation of the process is concerned, and with respect to the use of this formalism by clinicians.

Dr. Pauker joined the group in 1971, bringing to it a rare combination of expertise both in medicine and in computer science.

During the latter stages of our work on decision analysis, we began to see certain difficulties in using decision analysis as the sole basis for a system to deal with real problems of crisis medicine. After further definition of these difficulties, we were given a research grant from HISMA under which we explored these problems. From this exploration emerged a recognition of the need for a close cooperation with skilled computer scientists.

In order to promote a closer union between researchers in computer science and the workers in our group, we held last summer a week long conference on the problems of clinical decision-making and the relevance of advances in computer science to these problems. Attending the conference were five members of the M.I.T. computer science faculty (including Professor Marvin Minsky, the director, of the Artificial Intelligence Laboratory and Professor Edward Fredkin, Director of Project MAC) and the members of our group already mentioned. The major result of this conference was the recognition of the potential benefits to medicine of a strong computer science supported research program, and the complementary benefit to computer science of a close involvment in medicine.

At this meeting, we resolved to organize a research program which would bring together first rate computer scientists and clinicians in a coordinated study of the problems of clinical decision-making. This proposal and the work upon which it is based is the result of that collaboration.

Since that meeting, we have been actively pursuing research in this area. We have funded our activities through small amounts of money from various sources. Despite this limitation of resources, however, we are proceeding at a rapid rate. In addition to the research discussed in this proposal, we are attracting graduate students in computer science. Five graduate students are already working with us, and we would have more if more funds were available.

Professor Gorry has joined the faculty of the Electrical Engineering Department at M.I.T. and is working at Project MAC. Professor Sussman of the Artificial Intelligence Laboratory is taking an active role in our research efforts, and other faculty, notably Professors Fredkin and Minsky are advising us and encouraging our efforts. Most notably, Dr. Schwartz will be a Visiting Professor at M.I.T. next year where he can devoted increased energy to the research program.

All this causes us to be very optimistic about our ability to mount an excellent program of research and education in computer science and medicine. The critical problem now is not the people or the ideas, but simply that we lack funds. Because our work seems so well in line with the intention of the Computer Laboratory Program, we hope to obtain the needed funds from that program.

Previous Research on Clinical Decision-Making by Computer

Broadly speaking, work on computer-aided clinical decision-making falls into two categories. In the first category are efforts to develop computer-based mechanisms for assuring orderly and complete acquisition of data concerning the patient. Examples of such efforts are Weed's problem-oriented approach {6} and work in history-taking, physical examination , and laboratory testing procedures (See, for example {7}.) It is believed that with improvements in the data acquisition and data structuring processes will come improvements in either the effectiveness or the efficiency of the clinical decision-making process, and in general, this belief seems well-founded.

In the second category fall all the efforts which are directed at developing computer realizations of procedures for making diagnostic and/or management decisions. In general, activities of this type have paid less attention to the orderly acquisition of facts than to the problems of interpreting the facts as presented. Within this category, however, a further division of efforts can be made. This division is based on the view which the researchers take of the decision-making procedures they are developing -- whether these are thought to be <u>descriptive</u> or <u>normative</u>. In the former case, the researchers have attempted to codify the way in which experts actually make diagnostic or therapeutic decisions. In most cases, the determination of exactly how an expert behaves has been rather <u>ad hoc</u>, involving a mix of introspection, interview, and various forms of observation. Some notable successes have been achieved in this way. {8} (Here we are measuring <u>success</u> in terms of providing <u>distributable</u> expertise about some problem domain.)

Those workers with a more normative bent have emphasized the development of models and procedures for decision-making which are thought (under certain assumptions) to be the basis for <u>optimal</u> decisions. In almost all cases, the assumptions are met only loosely, and no real claim of optimality can be made. Still, the general flavor of the work suggests that computers <u>ought</u> to make decisions in this way, without regard to the way in which humans make the same decisions. The more normative approach has also yielded success in certain areas (e.g, {3}, {5}, and {9})

Although work in both of these categories has shown considerable promise, and research continues actively on both approaches, no program has been produced which can cope with the real complexities of the clinical situation, e.g. time dependent changes in disease, multiple disease in the same patient, and a variety of patient specific factors which have an influence on both diagnostic and management strategies.

> ***We believe that these approaches and the techniques which they have produced will enter into an expert system in an important way. We do not believe, however, that either of these approaches, as currently employed, can be the basis of the kind of administrative and integrative structure required in such an expert system.***

For this reason, we want to explore in some detail the methodological limitations of the approaches which have been used to date. It should be remembered that our criticisms of these approaches are in the context of trying to provide an overall framework for clinical decision-making.

1) Flow Charting

The 'descriptive' approach is to construct a flow chart to represent the way in which a particular problem is to be handled (e.g., {7}, {8}). As was noted above, the manner in which the flow chart is obtained is usually ad hoc. Sometimes the flow chart represents the opinion of an expert as to the process he believes he uses. In other cases, it is based on a mixture of introspection and more formal modeling of aspects of physiology or pathophysiology. In any event, the resulting flow chart is an encoding of a decision procedure which is deemed to be a good one to follow in the particular clinical area in question.

There are two major difficulties with this approach insofar as complex clinical problems are concerned. First , a rigid definition of the logic to be used in a given situation may be impossibly cumbersome if it attempts to account for time dependencies , multiple interacting problems, patient specific constraints, etc. Even if such flow charts can be constructed for subproblems of a clinical problem, the decision as to how and when they should be combined, modified, and applied to a given situation remains. The representation of knowledge in flow charts makes this latter decision exceedingly difficult. Medical knowledge about a given clinical situation is implicit, not explicit in a decision flow chart. Because the reasons for a particular branching are not available to the program, in general it cannot make even simple deductions about them. Thus, unless the clinical situation matches exactly a series of branches in the flow chart, the program is helpless, because its lack of underlying knowledge prevents it from adjusting its approach to a non-standard problem.

Further, with this kind of structure, a user cannot inquire about the basis for a decision or suggestion from the program. And, an expert cannot add new knowledge to the program except through a laborious search through the programs or frames of the flow chart to ascertain what the program already knows a given subject, and how the new knowledge should be related to it.

2) Decision Analysis:

Another approach to the problem of computer-aided decision-making is to give a program an explicit description of the relations between findings and diseases and between actions and outcomes. Then one can incorporate an inference procedure into the program for sequentially deducing the path it should take with respect to a given problem. This approach is the basis for the decision analysis program we built for acute renal failure ({3} and {5}), and has been used by others in different contexts. (e.g., {9})

By explicitly recognizing the uncertainty in the relationships and by generating a decision tree for each new situation, a decision analysis program for balancing costs and benefits can deal with the equivalent of a very large number of flow charts.

This work has demonstrated that decision analysis is a very powerful approach to problems of balancing risks and benefits in the clinical context.

With this approach, however, there are limitations which pose very serious problems when real-world complexities are introduced. Our current methods for the explicit description of the probabilistic relationships, the courses of diseases, action-consequence relationships, etc. are very rigid and to a large degree, artificial, and although these forms of description are well-suited for the decision analysis algorithm, they are very cumbersome for the expression of medical facts in medical terms. Thus, a time-consuming and error-prone process must be undertaken to translate descriptive statements (made by experts, for example) into material which the program can use correctly.

A second problem is that it is very difficult to give procedural advice to a program based solely on decision analysis. For example, an expert might want to suggest a logical procedure (perhaps a "flow-chart") by which a specific situation can be efficiently and effectively handled. He may have processed (in some way) all the uncertainties, risks, and benefits associated with the situation, and he knows that the procedure is useful. He cannot, however, add the procedure to the program directly. The options are either to reprogram the system or to determine some parameters which, when used by the decision analysis program, cause it to do the "right" thing. Both alternatives are unsatisfactory if much knowledge is to be added to the program.

Finally, to the extent that explicit descriptions of diseases, etc., are formulated in terms of probabilities, the knowledge of the program is basically a mass of numbers, and the explanation of decisions or suggestions made by the program will be very difficult for an expert (and more so for the average user) to understand. Concepts and language naturally employed by the expert to express his knowledge have to be converted to a set of numbers which when coupled with some decision produced the same results.

> *To summarize, neither the flow chart approach nor decision analysis can be the basis for a program which deals with complex clinical diagnosis and management problems. Both approaches have value in certain circumstances and should be used as appropriate, but new techniques are required for a program to be able to deal with the full range of complexities which arise in serious clinical situations. Advances are also required if it is to be possible for an expert to interact with a program in such a way that the program can assimilate the expert's knowledge, and for a user of that program to be able to have natural and direct access to that portion of the knowledge which is most relevant to the clinical problem he is considering.*

The need for these innovations is underscored by the diversity of knowledge which experts used. They use descriptive, causal, procedural, and administrative knowledge along with common sense. It seems apparent that current formalisms are suited for only one or two types of knowledge, and that a new framework for organizing and using these diverse kinds of knowledge is required. More recent work, such as that of the Rutgers Special Research Resource on Computers in Biomedicine is directed to the solution of some of these problems. We hope that the proposed Laboratory would establish close relationships with such activities.

The Relevance of Advances in Computer Science

Advances in computer technology, including dramatic increases in information storage capacity and the development of remote access capabilities in the form of time-sharing systems, suggest the possibility mentioned above, that computers would serve as a repository for medical expertise and as a means for disseminating that expertise to points of need within the population. If such 'knowledge-based' systems could be built to serve as consultants for clinical problems, they could be replicated (either in fact, or effectively through multiple remote access to one system) as needed.

Unfortunately, this computer power alone is not enough to carry us to our goal. As we noted in the introduction, the major impediment to progress is our lack of understanding of the processes of clinical cognition. Therefore, advances in computer programming and technology, alone, will not solve the problems. It is important, however, to recognize the role which advanced computer science and technology play in research such as that being proposed here.

It is an unfortunate fact that although advances in computer science and technology cannot solve the problems, deficiencies in either can pose a serious hindrance to progress. Until recently, various attempts to formulate behavioral theories of complex processes would have suffered from a serious lack in the existing technology, the technology which had to be the testing ground for these theories. As a result, the development of theories of intelligence in certain domains was retarded.

In recent years, there has emerged from research in computer science a new 'technology' for representing some kinds of knowledge in computer systems. This capability is relatively new, dating from the late 1960's, and we believe that its availability will greatly ameliorate the problems of formulating and testing cognitive theories. This in turn will have a very beneficial effect on research into clinical decision-making. We are not claiming that there are no technological problems in our path; on the contrary, there are many. It is our opinion, however, that this new technology permits us to begin to explore new forms of procedures which simulate aspects of clinical cognition.

The advances and ideas to which we are referring are concerned with new techniques for programming computers and new techniques for representing knowledge and meanings in programs. In the old style of making 'computer models', things were very <u>rigid</u>. In the new style, it is much easier to include knowledge about how contingencies and side conditions affect, not only the states of the models, but especially how the models are to be applied in various situations. (Later we will describe some of our ongoing research in applying some of these ideas to the problem of digitalis/digoxin administration.)

In the new style, communication between programs is more flexible and direct. Some kinds of knowledge can be represented as procedures, able to intervene actively in the control of other programs when specified 'patterns' arise in the other programs' operations.

<u>Goal-Directed Programming Languages</u>

Rather than being organized as a step-by-step sequence of actions to be performed, specified in advance by the programmer, programs in these programming languages are controlled by the activation of certain statements called <u>goals</u>. When a goal is activated, the system retrieves from a data base of knowledge statements those that match the 'pattern' of the goal. (A <u>pattern</u> is a description of a state of affairs in a model, or an encoding of some fact about the world, etc.) These retrieved statements then serve as <u>advice</u> about what should be done to achieve the goal; they may dictate that a certain program be run, that the goal be replaced by one or more subgoals, or that certain priorities be re-arranged., and then control be returned to an earlier, superior goal system.

Understanding Natural Language

For twenty years, the public has been titillated by promises that computers would understand natural language and even translate from on language to another. A justifiable skepticism has resulted from such promises. Although progress in the theory of 'syntax', both formal and informal was steady, this progress did not lead to the anticipated improvement in the computer's ability to handle language. The trouble, of course, is that syntax is not enough. A deeper understanding of the semantics of language was required. Only in the late 1960's with the work of such people as Winograd, Woods and currently Martin, were the earlier skirmishes with the problems of syntax and semantics sharpened into serious attacks on the problems of the meaning of language. (See, for example, {10}.) Thus although real problems remain to be solved, there is now justifiable optimism that a natural and direct interface between a user and a knowledge-based system can be built.

We want to underscore the importance of research on natural language to the kind of work we are currently doing, and to the proposed work of the Computer Laboratory. Of course, there is the obvious advantage of having a natural language interface with a program which contains clinical knowledge about some domain. Such an interface will permit the direct involvment of various experts (some not actively involved with the research of the Laboratory) with the program. This involvment will provide invaluable feedback with respect to the 'facts' in the program and with respect to the theories upon which the program is based.

A second benefit, perhaps, is less obvious. It has become clear that in large part the major impediment to progress in natural language research has been in semantics rather than syntax. The recent progress has built on new and better schemes for representing meanings. Further, as this research progresses, these representational schemes will be further developed and refined.

Even a cursory study of the kinds of knowledge employed by experts in solving clinical problems shows how much use is made of conceptual frameworks which at present are receiving increasing attention in language research. Such concepts as time, causality, change, etc. require deep analysis if machine representations of their meanings are to be found. The central role that such concepts play in medical knowledge means that progress by natural language researchers will almost certainly benefit our research directly. In fact much of our current thinking about representation of medical knowledge is strongly influenced by our colleagues (e.g. Martin) who are working on English.

Recognition and Analysis of Conflicting Goals

In many problem-solving applications, the recognition of conflicting goals is an important problem. Further, once these conflicts are recognized, it is important to have some means for resolving them. In earlier problem-solving programs, the recognition of goal conflict was generally difficult, because the goal structure of a program was implicit in the program itself. As we noted, the use of goal-directed programming languages lessens this problem considerably.

The analysis of conflicting goals, although still a significant problem, is also an area where improvements have been made. In the past, conflict between goals was handled by very crude strategies: either the goals were assigned simple priorities, or a trial-and-error search procedure would be tried first on one goal and then on the other in the hope that both would be achieved in some attempt.

Only recently have programs been developed which monitor their own performance sufficiently well to recognize and describe conflicts as they occur.Such monitoring is made possible in large part by the use of the goal oriented languages mentioned above to make the intention of a program more clear. (See, for example, [11]). Once in the open, problems of conflict can be faced (perhaps by special purpose programs) instead of being hidden in the rather arbitrary control structures of conventional programming systems.

Although we cannot say with any certainty exactly what processes would be needed for a computer simulation of the clinical cognitive process, it seems certain the performance monitoring and the analysis of conflicting goals would play important roles. Therefore advances from computer science research in this area are undoubtedly important for our proposed research efforts.

The Role of Computer Science Methodology

Perhaps the most important contribution which computer science research can make to the activities of the proposed laboratory is methodological in nature. The major reason that cognitive psychology has made relatively little progress with respect to understanding behaviors as complex as that involved in clinical decision-making is because there was a serious shortage of ways to describe the more procedural aspects of that behavior. As has been argued in [12]:

> *"The community of ideas in the area of computer science makes a real change in the range of available concepts. Before this, we had too feeble a family of concepts to support effective theories of intelligence, learning, and development. Neither the finite-state and stimulus-response catalogs of the Behaviorists, the hydraulic and economic analogies of the Freudians, or the holistic insights of the Gestaltists supplied enough technical ingredients to develop*

> *such an intricate subject. It needs a substrate of debugged theories and solutions to related but simpler problems. Computer science brought with it a flood of such ideas, well defined and experimentally implemented, for thinking about thinking; only a fraction of them have distinguishable representations in traditional psychology.*

It is this rich set of ideas which we plan to exploit in the description and analysis of clinical cognition. From this effort will come a new theory of the behavior of clinical experts and new concepts for the realization of this behavior in a computer.

Research Plan

Introduction

In order to provide a context for our discussion of the research plan for the Laboratory, we want to re-iterate our goals, and to relate these goals to our perceptions of the needs of the health care system.

> *We propose that the major activity of the Laboratory will be the use of the computer and advanced computer science methodology in the study of clinical decision-making. From the activities of the Laboratory will come two major results: 1) a deeper and better-articulated theory of expert clinical cognition, and 2) mechanisms for realizing the concepts of the new theory in computer programs for clinical decision-making.*

The reasoning underlying the organization of the Laboratory around these themes is as follows. We start from the premise that there is a need for distributable expertise concerning a number of clinical problems. Our particular interest is in the domain of serious medical problems, problems which are often potentially life-threatening. If we can make progress in understanding the way in which serious and complex problems should be dealt with by a clinician, and hence by a computer, we will be able to develop new technology of considerably improved flexibility and power which will be applicable across a broad range of medical decision-making applications. It can be anticipated, for example, that these advances will have an impact on the ability of the practicing physician to deal with complex or serious medical problems, placing the consultant as near as the nearest console. Such expertise should make far more effective the performance of allied health personnel, such as nurse practitioners and MEDEX personnel. In remote rural areas, for example, the availability of expert consultation should make it possible for allied professionals to deal competently with problems more serious than they otherwise could care for. In addition, the computer should be able to serve an important triage function, assisting the non-physician in his decisions concerning referral - in effect telling him when he should transfer the patient to a physician for care.

At present, however, the techniques for providing computer-based consultation are limited in application and remain generally incompatible with one another because no mechanisms for organizing and integrating them in a more general clinical context.

> *It is this lack of integrative mechanisms which is one of the principal impediments to the realization of the full potential of the computer in health care delivery.*

Our goal is to undertake the research which will produce these integrative mechanisms, and to do this, we have turned to the study of the behavior of clinical experts, because these experts have demonstrated their abilities to combine various approaches into a coherent strategy suitable to a given situation. We should begin by understanding how they achieve their performance. Recent advances in computer science provide us with new building blocks from which we can construct a better theory of clinical cognition. This theory will be developed through extensive use of computers and computer programs as a medium for <u>expressing</u> the theory, and as the means by which the theory is <u>tested</u>.

Below we will outline a set of research projects which we believe have the proper orientation to yield major progress toward the understanding we are seeking. As our work progresses, of course, new paths will become apparent, and our ability to define problems more sharply will increase.

In what follows, we have listed the <u>principal</u> participants in each project. Each group of principal participants contains computer scientists and clinicians, and the activities of the groups are fully collaborative. In a real sense, everyone mentioned in any project has an active interest in all the projects, but we thought it might be of some interest to the readers of this proposal to know who currently plays a major role in each project.

Taking The Present Illness

Principals
Professor G. Anthony Gorry
Peter B. Miller
Dr. Stephen G. Pauker
Dr. William B. Schwartz

'Will, as Schwartz has suggested, the computing science "largely replace the intellectual functions of the physicians?". I think not. The subtle process of the patient-physician interaction and the input we receive from this interaction has not yet been reduced to precise mathematical terms. Attempt as we will to analyze this subtle process, it appears that (despite) our best efforts to penetrate (it), this mystery will elude us for some time.'

(Warren Glaser, Professor of Medicine, in a comment on a forthcoming article on computer-aided diagnosis)

The sentiment expressed in this quotation is shared by many physicians. Those who have thought carefully about the interaction between a patient and a physician realize the complexity of the behavior involved. When a physician is confronted with a patient with one or more presenting problems, he enters into a mode of data acquisition and problem solving known as 'taking the present illness'. This activity is one in which virtually all clinicians participate every day. When we try to understand this process in detail, however, we find that it assumes a very complex and often subtle character. In fact, virtually all the problems of clinical cognition arise in this context. The process is like a puzzle for which some of the pieces can be rather easily found and described, but for which others remain quite vague and apparently ill-formed, while some appear to be missing entirely. The question of interest here is to what extent can we identify the pieces of that puzzle and put them together to form some coherent picture.

On the other hand, if a machine is to understand the process of clinical problem-solving, it must understand the taking of the present illness, because it is this process which provides much of the underpinning of the rest of the decision-making activities. Therefore, a deep understanding of the behavior of the clinician in this setting would provide a great deal of knowledge about how to support clinical decision-making. Additionally, we chose to begin work on the present illness because it represents the initial point of contact between patient and doctor, and because of the richness it presents with respect to cognitive processes and the integrative demands it places on the clinician. Further, it has the advantage that issues of risk and benefit such as those we addressed in our work on decision analysis can be ignored. Later, as our understanding increases, we can move the boundaries our our work to include these issues as well.

Preliminary Work

Our activities began with the analysis of protocols, tape-recorded records of the verbalized problem-solving behaviors of clinicians. These protocols are augmented by in-depth questioning of the clinicians regarding their approach to specific clinical problems and by criticisms offered by these same clinicians of preliminary computer realizations of our hypotheses concerning their cognitive processes. The purpose of all these efforts is to gain a deeper understanding of the way in which clinicians actually deal with the complexities of the clinical environment.

We have developed our hypotheses to the point where it has been possible to implement a rudimentary computer simulation of the process of taking a present illness. Though very detailed studies of the problem solving behavior of that program, we have gained new insights into the process. The use of the computer as the medium for the expression of the theory has aided enormously the advancement of that theory. This close man-machine exploration of the behavior of the simulation of the theory will be a key aspect of our research "style". Of course, this style has the additional benefit that when a satisfactory theory has been developed, a program which takes an excellent present illness for the given problem domain will also be available.

A further aspect of our style has been our emphasis on a "complete" examination of the issues involved in taking a present illness for a single complaint (in this case, edema). By forcing ourselves to consider even "minor" differences between the behavior of the program and the behavior of the clinician as problems for investigation, we have considerably sharpened our understanding of the process the doctor uses.

Now we want to present our first, rather rudimentary understanding of the problems and processes associated with the present illness. Then we will describe our first theory and the computer realization of that theory. Finally we will discuss our research plans for this project.

Observations of the Present Illness

The physician, when taking the present illness, asks the age and the sex of the patient, and elicits a chief complaint. The latter is the problem which caused the patient to seek medical attention, but it will often be closely followed by mention of other problems the patient has. In fact, one interesting problem which is currently of concern to us is how a clinician links several presenting problems together. For simplicity of discussion, however, we will assume that the patient presents with a single chief complaint.

The response of the physician to the chief complaint will vary in details, but the principal thrust of it will invariably be at elucidating and refining the description of the complaint as given by the patient. For example, if the patient's chief complaint is 'swelling of the face', the physician's questions generally will explore the duration of the swelling, its specific location (e.g. around the eyes), the symmetry of the swelling (is it only on one side of the face?), etc.

The characterization of the presenting complaint is important because it is this characterization which along with the age and sex of the patient gives the clinician his initial framework or <u>context</u> within which to work.

> ***The rapid selection of a context is vital for the clinician. The clinician is about to hear a reasonably large amount of information from the patient, and if he is to be able to organize that information and to deal with it effectively, he must have a framework into which it can be fitted. Because of the breadth and diversity of medical problems and the scope of knowledge concerning these problems, a failure to focus attention and to narrow drastically the domain under consideration will prevent the clinician from <u>understanding</u> what he will be told.***

Note that this is the reason physicians require the age and sex of the patient at the outset of the history; because these facts, in conjunction with the chief complaint provide a great deal of <u>focus</u> for what follows. Consider the difference in your reaction to the chief complaint of 'severe, progressive weakness' in the case of an 80 year old man, and that of a 13 year old girl.

Therefore the initial goal of the physician in taking the present illness is to get an <u>adequate</u> description of the chief complaint of the patient. What constitutes an adequate description, however, is determined by another fundamental goal, namely that of gaining a framework within which to understand the information which will be forthcoming from the patient.

In some cases, fragments of this investigation will appear to be a rote recitation of a standard sequence of questions (e.g. in the case of abdominal pain, 'Is the pain made worse by lying down?', 'Is it made worse by eating?', 'Is it made better by eating?', etc.). Other fragments will be strongly influenced by the responses of the patient. For example, if the swelling of the face is periorbital and symmetric, the physician might want to know whether it appears in the morning and disappears during the day. If the answer is yes, then he might well transfer his attention to an investigation of possible pedal edema. On the other hand, if the swelling is in one cheek and is painful, the investigation might switch to questions of recent dental work on the patient.

Clearly, then, the path of the investigation of the chief complaint taken by the physician is in part a function of the responses given by the patient to the former's questions. This path is equally well a function of the clinical knowledge of the physician. Only a doctor who recognizes the periorbital edema described above as very likely the result of renal disease (specifically acute glomerulonephritis or less often, nephrotic syndrome) would follow the path suggested. So underlying the observable behavior of the physician is a knowledge base, the use of which is only implicit in the process of investigation.

That the investigation of the chief complaint follows a path determined by both the medical (and other) knowledge of the clinician and the responses and descriptions given by the patient is apparent to anyone who has looked at the present illness in even the most cursory manner. Thus it is non-controversial that these two factors are pieces of our puzzle. What remains unclear is how these pieces interlock in any given situation.

The exploration of the chief complaint generally results in a much sharper characterization of it than originally offered by the patient, although usually only certain additional features of the complaint have been elicited, i.e., the exploration of the complaint has been stopped short of exhausting all the properties which this problem might conceivably have. This of course raises the possibility that some aspect of the patient's problem has been overlooked, and the need for further investigation may arise in later in the session.

The characterization of the chief complaint as elaborated by this process can prompt a number of different behaviors on the part of the physician. In certain cases, the description of the complaint suggests little to him, and so he may simply encourage the patient to volunteer more information ('Have you had any other difficulties lately?') or he may begin a 'review of systems' type of investigation of the system involved in the patient's problem.

If the latter approach is used, however, it will seldom persist as the basic <u>modus operandi</u>, because it is too passive for use in taking the present illness, and it is used here only as a temporizing measure. As soon as it yields some additional information, the physician will assume a more aggressive stance with respect to information gathering.

The purpose of this excursion into the review of systems is the same as that underlying the original attempt to refine the characterization of the chief complaint, namely to get just enough information to glean a good suggestion of a <u>context</u> for further discussion of the patient's problems.

The initial context chosen will of course be further refined as the present illness is taken. It may be an organ system (in the sense that the chief complaint is strongly suggestive of a problem with that organ system); it may be much more specific in that the chief complaint might suggest a specific disease. (Of course, there may be more than one disease or organ system suggested.) In any event, the extent to which the clinician pursues the characterization of the chief complaint depends on the search for an appropriate context and the potential availability of contexts which are quite specific. For example, the facial edema described by the patient above would be pursued to establish its specific location and temporal pattern because of the specificity of the renal disease context which would result if the appropriate characterization could be achieved.

At its most macroscopic level, the taking of the present illness can be described as the clinician moving from context to context with occasional returns to previously-invoked contexts. At each context, the activities of the present illness can be thought of as being under the control of that context. By this we mean that the questioning of the patient is directed at either the confirmation of details associated with the context (such as asking about pedal edema because it is generally found when periorbital edema is present) or at the selection of a more 'specific' context (as when the clinician asks a patient with exertional dyspnea whether he has paroxysmal nocturnal dyspnea in order to choose between the contexts of lung disease or heart disease).

<u>Present Illness or Diagnosis?</u>

Before we continue our discussion, we want to comment on the role which <u>diagnosis</u> plays in the present illness. Clearly, the present illness is 'driven' by the desire to establish an understanding of the patient's problems and their interrelations with one another; hence the clinician is seeking a diagnosis which is suitable as a basis for management decisions. There is a very real sense, however, in which the present illness is more than diagnostic process as the latter is conventionally construed.

Normally we think of <u>a diagnosis</u> as an inference about the state of the patient which is based on his signs and symptoms, and we call the activities associated with the collection of information (identification of signs and symptoms) the <u>diagnostic process</u>. We have noted that the taking of the present illness is also an information gathering activity, but it is directed as much toward the problem of ascertaining <u>what the facts are</u> as it is toward the problem of <u>what the facts mean</u>.

Although we admit that there is a level at which one can view the present illness as part of the diagnostic process and the process of diagnosis as an integral part of the taking of the present illness, we feel that the distinction we have made has some merit. It helps expand our view of the problems of clinical cognition.

For example, when we think only of 'the diagnostic process' we tend to think of such questions as 'What inferences can you draw concerning a 28 year old man with dyspnea and orthopnea who had an attack of acute rheumatic fever when he was 15, and.. . etc.' We tend to view the problem as understanding the meaning of a constellation of findings <u>as given</u>. We assume that the patient indeed does have dyspnea and orthopnea and that the attack of rheumatic fever actually took place. In taking the present illness, however, the clinician often is not given these facts, but must 'dig them out', and even then he may be left with significant doubts concerning the facts themselves. It is this additional aspect of establishing and characterizing the facts and assessing their reliability which we are emphasizing in our rather arbitrary distinction between the process of diagnosis and that of taking the present illness.

Now that we have made the point that the two activities of <u>establishing</u> the facts and <u>interpreting</u> the facts are central to clinical cognition, we will now explore some of the ways in which these two activities interact, and we will drop our distinction between taking the present illness and working toward a diagnosis.

<u>Prerequisites for Clinical Cognition</u>

Although many of the details of the processes employed by the clinician in taking a present illness or in proceeding to a further diagnosis are still obscure, it is possible to identify some major aspects of the general cognitive process. We can do this by analyzing the <u>task environment</u> of clinical medicine. A physician who is well adapted to that environment will <u>necessarily</u> possess cognitive processes for dealing with each of the major demands placed upon him by the environment. Although we may not be able at present to give much detail concerning these processes, we will have made a first step by recognizing the necessity of their existence. (In the following discussion, we make use of some terms borrowed from Minsky [13].)

1) <u>Expectation and Focusing</u>

The first problem that a clinician faces when he is dealing with a patient is that both the number of disease states and the number of possible findings which may have some relevance are extremely large. This means that the clinician faces a search through a potentially bewildering maze of possibilities. Because his cognitive capacities are limited (especially with respect to the number of 'simultaneous' paths he can explore), he must use the facts as presented to <u>drastically</u> reduce the number of possibilities which he will consider in any detail.

As we noted in our brief discussion of the present illness; this rapid focusing serves the principal purpose of providing the clinician with a context for his further problem solving activities. In our studies of expert clinical decision-making, we have been struck by the rapidity with which experts achieve such a framework. When they are presented with only a few (two or three) facts, experts almost always have one or two working hypotheses. It may very well be that the hypothesis first chosen will later be discarded. Our point is not that this first choice is an <u>accurate</u> or <u>optimal</u> one. It is a <u>good</u> working hypothesis, however, in that it brings important <u>structure</u> to the problem.

Because the stimuli for this focusing are the presenting signs and symptoms of the patient, it is reasonable to infer that the expert remembers patterns of findings which 'point to' good working hypotheses or contexts for those findings. Our current speculation is that these patterns contain relatively little detail, and they serve only as a first rough cut at the problem of classifying the patient. This speculation is based primarily on the experts' descriptions of the patterns they are using and on the rapidity with which this focusing takes place.

When a context has been selected, the clinician appears to match the findings of the patient against a more detailed description of the prototypical pattern of findings associated with the context. For example, 'shortness of breath in a 50 year old man' immediately suggests the contexts 'heart disease' and 'lung disease'. (Notice in fact how focused these contexts are relative to the total number of disease states which could be presented by the patient.) Most clinicians would proceed immediately to the characterization of the shortness of breath in order to focus on either heart disease or lung disease.

This attempt to match the presenting findings or the chief complaint to a more detailed pattern for a context is typical of the activities which underlie much of the present illness. For example, consider the presenting problem of periorbital edema. It immediately suggests (among a few other things) acute post-streptococcal glomerulonephritis. A renal expert would very likely move directly to a series of very detailed questions concerning the temporal pattern of the edema. The context of AGN has already been 'suggested'; the detailed examination of the characteristics of the edema will determine whether this context will govern the succeeding questions of the clinician.

2) Elaboration

Once a context has been chosen, the clinician faces the problem of confirming his choice. This confirmation requires two steps: first, he must convince himself that the rest of the signs and symptoms presented by the patient conform to his understanding of the disease state or the physiological state represented by the context, and second, he must assure himself that these findings are not better associated with one another in some other context.

One of the fundamental principles which we have observed in out studies is that experts use the principle of parsimony. The expectation that all the patient's findings are related to the same problem is strong in the clinician's mind. He yields this idea only grudgingly. In our discussion below, we will see examples of the major role this idea plays.

The process of elaboration is very complex, involving several distinct, but interacting activities.